Questions and Answers

Chapter 1

Review Questions

1. You've just received a new product that you ordered. Your boss walks by with someone from another division who's considering the same product. He wants to know whether the product has a particular feature. How would you go about finding out for sure?

   The best place to start is probably to log on to the vendor's Web site and search for the name of the feature. If that's not effective, you might try skimming the documentation for a mention of that feature or a similar feature. If you don't find it there, you could call technical support.

2. You're evaluating three different products with similar features and similar costs. You're concerned about making the right decision. How should you determine which is the best product?

   The information you need at this point is not available in the spec sheets or from the sales reps. You need to see what the experience of other users has been. The best way to do this is to search the Internet newsgroups for feedback.

3. You've been on vacation for a week, and you get in early your first day back to read your e-mail and get settled. As you're walking in, however, a user stops you and says that she can't log on. After a quick check at her computer, you determine that you can't log on either. What should you do to figure out what is going on?

   The first step is to try to log on at your workstation. That will determine whether it's a problem with the user's computer or with the entire network. Assuming that it's not a network-wide

problem, check the network activity log to determine what changes were made while you were gone.

4. You've done the necessary analysis for a new server. You've taken into consideration the needs of the organization and have quantified the issues around the need. However, management doesn't seem pleased with the solution. What might have you forgotten?

   Perception is reality. Even though you may be able to quantify the cost of downtime and the expenses that you've included in your solution to address risks, management may have a different perception of what is risky and what isn't.

5. Your company is planning a phase of rapid growth. You're trying to determine the best course of action for purchasing servers. What options should you consider?

   There are a few approaches that you may need to pursue in parallel. The first is trying to purchase a server that will accommodate the projected load. Another approach is to set up a small clustered environment and then add servers to the cluster as the load increases. Finally, you could consider adding a single server now to handle all services and then slowly shift services to new servers as the load increases.

6. You're getting ready to add a new server. What method might you use to choose the operating system to run on that new server?

   You can choose it based on what you know or based on what is popular. However, the best technique would involve factoring in the benefits and weaknesses of each operating system.

7. Which will perform better, hierarchical PCI buses or peer PCI buses?

Hierarchical PCI buses use a root PCI bus for all transfers, while peer PCI buses allow for concurrent communications on different buses. Because of this, peer PCI will have a performance advantage over hierarchical PCI bus designs.

8. When is hot-swap PCI important?

   When availability is paramount. A server with hot-swap PCI slots is generally more expensive, but it does allow you to remove and replace some PCI cards without powering down the server. This can be a great advantage in 24-hour operations.

Chapter 2

Practice: Size a UPS System for a Small Environment

1. What is the total number of volt-amps that are necessary?

   A 1680 VA UPS is necessary to carry the load.

2. If you require 30 minutes of run time, which UPS is necessary?

   To run for at least 30 minutes, the 3 KVA UPS is necessary.

3. If you require 60 minutes of run time, which UPS is necessary?

   To run for 60 minutes, the 6 KVA UPS is necessary.

Chapter 2

Practice: Size a UPS System for a Larger Environment

1. What is the total number of volt-amps that are necessary?

   A 3960 VA UPS is necessary to carry the load.

2. If you require 30 minutes of run time, which UPS is necessary?

To run for at least 30 minutes, a 6 KVA UPS is necessary.

3. If you require 60 minutes of run time, which UPS is necessary?

   To run for 60 minutes, two 6 KVA UPSs would be necessary.

Chapter 2

Review Questions

1. Why is climate control important? What impact does it have on performance or availability?

   Climate control ensures that components of the server are not damaged or prematurely aged by excessive heat. It also ensures that the contacts don't corrode due to excessive humidity, leading to long-term reliability problems. The impact of climate control is solely on availability or reliability. Controlling the climate into which the server is placed have no impact on performance.

2. Your organization's business offices are located on the Florida Keys, where hurricanes are a real concern. The organization wants to make sure that if a natural disaster strikes it will be able to get back up and running right away. What is the best disaster recovery solution?

   A hot backup site is the most effective solution because it offers the quickest recovery time in a natural disaster situation. Having exact duplicates of the systems means that all you need to do is load the systems from the latest backup.

3. Your organization is an application service provider (ASP), hosting applications for companies across the United States. One of the biggest concerns that customers express to the sales people is the potential for your systems to be down, preventing them from getting their work done. What power precautions should you take

to minimize the risk that a power failure could keep your systems from operating?

The obvious power protection involves having surge suppressors and UPSs. You would probably also need a generator because you may need to keep the systems up for several hours while the local utility resolves a problem.

4. Your organization has a computer room with a raised floor. Some people in the organization say you shouldn't be worried about flooding because the computer room is 12 inches above the floors in the rest of the building. Why isn't this true?

   Although it's true that the computer room floor is 12 inches above the rest of the building flooring, there are power cables running on the subfloor of the computer room, which is the same height as the rest of the building floors. Those power cables can short out and cause problems or fires if they come into contact with water.

5. You're installing a new server and you don't have any new circuits available to you, so you have to use an existing circuit. What should you do before plugging in the new server?

   Evaluate the load on the circuit. Remember that the circuit is connected with a circuit breaker or fuse that will shut down power to everything on that circuit if too much power is drawn at one time. If you're not sure how much power is on a circuit, have an electrician give you an idea; this can be done without disrupting the power to the existing devices.

6. Your organization has just hired a new director of information technology, and as a part of his initiation to the organization he wants to know what everyone does and has asked to see any materials that you have to show the processes you go through. You're responsible for monitoring backups. When you show him

the five failed backup reports over the last six months, he seems concerned. What might be concerning the new director of information technology?

The new director may be concerned that you're monitoring backups only for failures and not for successes and thus wouldn't notice if the process never finished and therefore didn't generate an error message. Monitoring for success, not just failure, is particularly important when the system that might be failing is not tested frequently.

7. You're installing a new server in a new rack, complete with a very heavy rack-mount UPS. Your boss likes the fact that the UPS has LEDs on it indicating load and battery availability and wants the UPS mounted at the top of the rack where everyone will be able to see these LEDs. Why is this a bad idea?

   Rack installation guidelines require that the heaviest items be placed at the bottom to help stabilize the rack. By placing the UPS on the top of the rack, you risk having the rack tip over and harm someone. If there's some doubt about the need to mount the UPS at the bottom of the rack, refer to the documentation provided by the rack manufacturer.

8. Your organization is in the pharmaceutical industry and must maintain records of clinical studies for several years. Although most clinical study information is archived on optical media, you're concerned about the ability to restore the study data from the regular tape backup and archive mechanisms. What consideration should you give to tape backup systems?

   You should consider the life of the medium that the drive uses. Although most server-class tape backup systems have long life spans, some don't. Because you might need to recover information from a tape, you should give special consideration to tape life.

Chapter 3

Review Questions

1. You're working for a medium-sized organization that needs a new terminal server. What kind of server should you select?

   When trying to select hardware for a terminal server, the priority should be squarely on processor and memory. A multiprocessor AMD Athlon system with a gigabyte or so of RAM would be a good start. Although SCSI would be preferred, you could probably get by with an IDE drive in a pinch. By using a multiprocessor AMD Athlon system and ample memory, you can get high performance without paying the premium for the fastest machine.

2. You're putting together a file and print server for an organization with a few hundred users. You're trying to determine how much cache to order for the processor. What guidelines should you use?

   The guidelines for selecting a cache size are based on the amount of memory that will be used and the relative dependence on the processor. Since memory isn't specified, the rule of thumb that file and print servers don't traditionally have large amounts of memory, and the knowledge that file and print servers are rarely processor limited, should lead you to choose a smaller processor cache size.

3. A soft memory error has been recorded on a server. What should be done?

   A soft memory error indicates that an ECC memory module detected and corrected a failed bit of memory. The memory module should be replaced as soon as is practical. Although ECC memory can correct single-bit errors, it cannot correct multiple-bit errors.

4. You're putting together a server to be a firewall. What kind of hard disk is appropriate?

   Since firewalls don't access their disks frequently, or in response to user requests, an IDE hard disk makes sense for this server. It's cheap and fast, and the multiuser performance implications don't apply.

5. Your organization has grown and has just built a huge new data center. It wants to standardize on a set of high-performance drives in one cabinet. What kind of bus is appropriate?

   Fibre Channel is appropriate for this kind of connectivity because the organization is obviously large (because of the new data center) and therefore has the potential need for very powerful servers. Also, because the drives are going to be kept in one area, the additional cable length that fiber optic cable provides may be necessary.

6. Your organization recently upgraded the backbone to support additional growth and has given you the task of selecting a new file and print server. You've settled on CPUs, memory, and hard disks. What kind of network connectivity should be provided?

   Given the size and growth of the organization in the question, a Gigabit Ethernet connection is probably warranted. Additionally, because the server will be the primary file and print server, a redundant, or fault-tolerant, NIC might be appropriate.

Chapter 4

Review Questions

1. What is the greatest risk when installing new firmware?

The greatest danger is that the power will go off during the programming phase of the firmware installation. This can render a card unusable.

2. What precautions should you take when installing new firmware?

   You should minimize the system configuration and test the power. You should also save the previous version of the software and take any additional precautions recommended in the README file.

3. When should firmware updates be performed?

   When the system is initially installed. After that, they should be performed only when you need to address problems or performance issues. Firmware updates are inherently risky and should be performed only when there's a specific reason.

4. What kind of a reboot should be performed after a firmware installation?

   Unless you're specifically directed otherwise in the firmware utility documentation or installation procedures, all reboots during a firmware reload should be cold reboots. That is, the power should be removed for a period of not less than five seconds and then reapplied.

5. What BIOS settings should never be changed?

   Those options pertaining to clock cycles, wait states, or memory speeds. These settings are set by the manufacturer and should not be changed unless technical support specifically directs you to do so.

6. Several options deal with video settings. How should the video settings be set up for a server?

In general, a server has little need for advanced video features and enhanced video performance. The server is better off conserving resources for operations other than video display. Video BIOS can be cached, but the memory aperture size should be set to its smallest size, and an IRQ should not be assigned for the video card.

7. You are installing Novell NetWare as your server operating system. How should the PnP/PCI settings be set?

   As of this writing, Novell NetWare is not Plug and Play compatible, so the BIOS should be set to indicate that a PnP operating system is not present. Configuration should be done manually via the BIOS and through PnP utilities provided by the manufacturers of the cards installed in the system.

8. You've added a new hard disk to a RAID controller and want it to be used as a hot spare. What must you do to accomplish this?

   You must flag the drive as a hot spare drive in the RAID controller's BIOS or in the RAID controller's configuration utility.

Chapter 5

Review Questions

1. What are some of the benefits of installing an operating system across the network?

   The primary benefit is that there is no chance of a media error. Additionally, the system will always try to locate the files at the original installation location when additional features are installed. Finally, network installations can be scaled to dozens of machines without creating or locating a large number of installation CD-ROMs.

2. Why do in-place upgrades take so long?

   In-place upgrades take a long time because of the need to back up the system and make changes slowly and deliberately. Sometimes they take longer because the server hardware is older and less compatible with the new operating system. When doing an in-place upgrade, you need to preserve the data that already exists on the system. In a new installation, there isn't any data to worry about backing up.

3. You're a busy professional who needs to maximize the effectiveness of your time. You're planning on loading a new server over the course of two days. What steps should be completed the first day?

   At the very least, you should establish the RAID arrays in the RAID controller BIOS. Because building the RAID array can be the most time-consuming process, you should start it and allow it to run. It might make sense to try to get to the drive formatting step before quitting for the day, because the RAID array build and the drive formatting could run simultaneously and be complete by the time you return in the morning.

4. What is the first step in installing most operating systems?

   The first step for most operating system installations is a basic hardware scan to make sure that the installation program will be able to complete and that the operating system will work on the hardware. This step generally finishes quite quickly.

5. What partitions are required for a Windows 2000 installation?

   A boot and system partition, although a single partition can be both the boot and system partition.

6. Why is the correct setting of the date, time, and time zone so important?

In today's environment of unified security systems such as Novell's NetWare Directory Services (NDS) and Microsoft's Active Directory, a great deal of synchronization of data occurs. To reduce the number of problems with data synchronization, it's important to keep the dates, times, and time zones correct and synchronized with one another.

7. What unique partition is used to boot a Novell server?

   A DOS partition is used to boot a Novell server. The server starts from a DOS partition of no less than 50 MB and continues booting by using the NetWare SYS volume.

8. What NetWare volume is required for a NetWare server?

   SYS, short for system. This volume contains the system and public utilities, in addition to various other components. It is required for any NetWare 5 installation.

9. Will a .DSK module work with NetWare 5?

   No, NetWare 5 requires the NetWare Peripheral Architecture that uses Host Adapter Modules (HAMs) and Custom Device Modules (CDMs). .DSK modules were used in prior versions of NetWare.

10. What kinds of partitions are required for a Linux installation?

    Linux requires a boot partition to boot from, a swap partition to use for virtual memory space, and a root partition for the system. You can omit the boot partition, in which case the system will boot from the root partition, but this configuration is not recommended. Additional partitions can be created but are not required.

11. What is a service pack?

A service pack is a collection of previously released patches that are assembled into a single package for installation on computers. Service packs are useful because they are extensively tested for potential interaction in various environments, whereas individual patches get less attention.

12. When should a patch be applied?

When there's a clear indication that the problem that the patch solves may occur or has occurred in the environment. In general, patches that fix problems that are unlikely to occur in the environment or that will not cause harm if they do occur should not be installed unless thoroughly tested.

13. A user name and password are used to determine what?

Identity. The identity of the user is assumed once a user name and password combination is received. Security or authorization is based on the access control lists to which the user belongs.

14. What is CRUD?

CRUD is an acronym for create, read, update, and delete. These are the basic accesses supported by all server operating systems except for Linux and other varieties of UNIX.

Chapter 6

Review Questions

1. Which message is generally used to initiate automated shutdown of a server?

The "low battery" message is generally the one used to start the unattended shutdown. This is because the "on battery" message may occur frequently.

2. How can a UPS be connected to a server for notification of events?

   Either by serial cable or via a TCP/IP-based network. Almost all UPSs support serial communications natively. The TCP/IP-based network option generally requires an add-on card or external device.

3. What might cause the input amperage to be higher than the output amperage in a UPS?

   A slight difference may be caused by the UPS electronics themselves, but most frequently this is due to the UPS charging its batteries.

4. What does an "emergency power-off" message in the event log mean?

   A user initiated an emergency power-off, via either direct keypad entry or a remote power-off switch. This message should not occur during normal operation. The emergency power-off option should be used only when there's potential for damage to the servers or injury to people.

5. Why is battery temperature important?

   Battery life is directly affected by the temperature of the battery. Batteries are heated when being charged as well as when being discharged. The charging of the batteries is regulated so that they don't overheat.

6. Drive cabinets generally have three connections to them. What are they?

   Power, drive bus, and monitoring. The need for power is obvious. The drive bus allows the server to communicate with the disk drives contained in the enclosure. The monitoring

connection allows the server to keep track of the condition of the cabinet itself, including the power supplies and fans.

7. In Linux, how can you determine the version of a driver?

   If you have the C source file, you can review the first few lines of the driver. Otherwise, the console log is the best place to look for the driver's version. Forcing the driver to reload, if it's a dynamically linked module, can generate console messages that might contain version information. These messages will appear at the bottom of the console window.

Chapter 7

Review Questions

1. What are the name and location of the Novell NetWare log file that records critical errors?

   SYS:SYSTEM\ABEND.LOG. This file contains the dump information created each time the server reboots due to an abnormal end.

2. What is the file that Linux records when an abnormal process termination occurs?

   Abnormal process terminations in Linux generate core dumps.

3. How does the procedure for resolving progressive errors differ from that for resolving critical errors?

   Critical errors require immediate resolution. The procedure for resolving critical errors is aggressive and focused. The procedure for resolving a progressive error involves a slow, calculated effort designed to identify the root cause and eliminate it.

4. What are the two different signs of a progressive error?

   Increasing severity and increasing frequency. Increasing severity of errors means that the errors themselves are getting more severe. Increasing frequency means that the same error is occurring more often.

5. Why is it important to review the log files of a server?

   To identify progressive errors and to manage the size of the log files.

6. What is the shutdown procedure for a Novell NetWare server?

   You shut down a NetWare server by issuing the DOWN and EXIT commands in sequence.

7. How can Linux be shut down?

   Linux can be shut down in a variety of ways, including the reboot, shutdown, and init commands.

8. What two settings should be established for unattended shutdowns to be completely effective?

   The UPS should be set to restart a few minutes after utility line power is restored. The server should be set to boot when the UPS resumes supplying power to it.

Chapter 8

Review Questions

1. What is performance monitoring?

   Performance monitoring is the act of collecting performance-related statistics so that decisions regarding performance and capacity planning can be made.

2. What are the four basic categories of counters?

   Memory, processor, disk, network.

3. What is SNMP?

   SNMP, short for Simple Network Management Protocol, is a protocol that can be used, among other things, to monitor the performance of computers running Linux, Windows 2000, NetWare, and network devices such as switches and routers.

4. Why is the amount of memory being used for a disk cache important?

   Disk caching reduces the number of times that the server must access the disk. In general, the more physical memory being used for disk caching, the better the overall performance will be. (However, beyond a certain point additional memory will not be used for cache and will not be helpful.)

5. What size disks are appropriate for database servers?

   Smaller disks are generally appropriate for database servers because the smaller capacity will require more disks to reach the necessary total capacity. Each additional disk in an array improves the maximum throughput. Database servers generally require more disk access for the same amount of disk space, so higher throughput is desirable.

6. Why shouldn't you use the monitoring tools on the server to determine whether the network is the bottleneck?

   Network performance monitoring shows only the traffic seen by the server. When switches and routers are used in the network environment, it's impossible for the server to show the complete traffic pattern on the network. Appropriate tools for network monitoring are an SNMP management console

capable of reading the routers and switches on the network and/or a packet sniffer.

7. What are non-unicast packets?

   Unicast packets are packets directed to a single destination. Non-unicast packets are packets destined for more than one machine. Typically these are packets that were broadcast from a machine to all other machines on the same network. Limiting non-unicast packets is a good idea, particularly when switches are present in the network.

8. What is the most common cause of inbound network errors?

   Collisions. Ethernet uses Carrier Sense Multiple Access/Collision Detection (CSMA/CD). This method of access will generate a small number of collisions that are read by the network adapter as inbound errors.

9. What mechanisms can you use to control the size of a performance log?

   To control the size of performance logs, select only the counters that are necessary for analysis, and change the frequency at which the performance data is captured to a longer interval.

10. What is a baseline?

    A baseline is a meaningful reference point for performance data. It serves as a basis for comparison with later data. Baselines are always run on the same server and should be run under the same initial production load that subsequent performance data will reflect.

11. Why should you review performance data in a specific order when identifying bottlenecks?

Some areas of performance can mask other areas. By analyzing and assessing the performance data in an orderly way, you can ensure that you're looking at the root cause of an issue.

Chapter 9

Review Questions

1. Why do we document a server's installation?

   To make it easier for us and others to do capacity planning, maintain the server, and provide support.

2. Why can't we document every detail?

   There are too many options to document every one. The goal of this type of documentation is to describe the framework for the installation, along with details that are difficult to determine or get. It is not the intention of documentation to take the place of interactive investigation of the operating system

3. Why do we make copies of invoices and purchase orders relating to the server?

   To ensure that we can provide proof of purchase if the vendor requires it for technical support, even after the accounting department has archived the information pertaining to the purchase.

Chapter 10

Review Questions

1. What is a thread?

A thread is the smallest unit of execution within an operating system. Threads, not applications, are scheduled to run within the operating system.

2. What is processor affinity?

   Processor affinity is an operating system setting that controls how the operating system schedules threads to be run. A higher processor affinity means a greater chance that a thread will be scheduled to run on the same processor as the last time.

3. What is thread binding?

   Thread binding is a processor affinity setting that requires the operating system to run a task on one of the processors. This setting can be used for applications that do not run well on multiple-processor systems and thus need to run on only one processor. The use of thread binding is generally not recommended because it doesn't allow the operating system to optimize itself.

4. What are the requirements for adding an additional processor to the system?

   You need operating system support, the same or similar processor stepping levels, and a processor slot that can have a processor inserted in it without interfering with any add-on cards. Not all operating systems support multiprocessor systems, and those that do sometimes require different installations or licenses. Maintaining similar processor stepping levels ensures that the timings within the system won't vary so much as to be a problem. A power supply capable of driving the additional processor is also required.

5. What are the three concerns you need to address when upgrading an existing processor to a faster one?

Memory speed, core voltages, and clock rates. As the processor speed increases, so must the speed of the memory. As processors get faster, they generally require less voltage to power their core, and higher voltages will cause damage. Clock rates are controlled by the motherboard, so if a faster processor is installed on a motherboard that won't support the higher speed, there won't be any speed improvements.

6. After you've physically installed the memory and verified it in the BIOS, what is left to do?

   Verify that the memory is detected by the operating system, or register the memory with the operating system. You may also need to resize the virtual memory files and partitions.

Chapter 11

Review Questions

1. What are the two types of cables used for ATA/IDE devices?

   ATA/IDE devices use either a 40-conductor cable with a 40-pin connector or an 80-conductor cable with a 40-pin connector. The hybrid 80/40 cable is required for IDE devices operating at Ultra/66 or higher speeds.

2. Why do ribbon cables present connector problems?

   Ribbon cables themselves don't cause problems. However, the pin block connector type that is frequently used with ribbon cables doesn't have an inherent method for orienting the connector the right way.

3. What three methods are used to identify how pin block connectors should be plugged in?

Pin 1 identification, blocked pins, and keyed connectors. Pin 1 identification is present on almost every cable. Blocked pins are rarely used any longer. Keyed connectors are often used with ribbon cables today.

4. How many types of connectors and cables are there for internal SCSI connections?

   There are two types of connector and two types of cable for internal SCSI connections: a 50-conductor cable with a pin block connector and a 68-conductor cable with a high-density D-shaped connector.

5. Which disk drive buses need to be terminated?

   Only SCSI buses require manual termination. Every bus has terminators; however, with ATA/IDE and Fibre Channel the termination is automatic.

6. What does the term "cable select" mean with respect to the ATA/IDE bus?

   Cable select is an older mechanism that identifies which device on an IDE bus is the master and which is the slave, based on the location of the device on the cable. However, it's rarely used anymore, as most ATA/IDE cables do not support cable select.

7. Which type of fiber optic cable is used most often with Fibre Channel?

   Multimode fiber optic cable. Multimode cables use lower power transmitters and have a maximum length of approximately 200 meters.

8. How many addressing modes does Fibre Channel support?

Two. Hardware addressing, in which internally burned IDs are used to reference the device, and physical addressing. Physical addressing identifies a device by its physical location on the Fibre Channel network. Hardware addressing is similar to the type of addressing used by networks.

9. What is the first thing you should do to address airflow issues?

   Check the cable routings and try to reduce their interference with airflow through the server. Drives should also be spaced out if possible.

10. What type of fan should be used first when an airflow issue is suspected?

    An auxiliary fan that pulls air out of a server (in the front) or that pushes air into a server. The goal is to move more air through the server, rather than moving the air in the server across more components.

11. Why are sleds used with hard disks?

    They are used to simplify the installation and removal of disks. Most sleds allow drives to be inserted and removed without removing power (called hot swapping). Once a drive is installed on a sled, it can be inserted quickly.

12. What is the primary disadvantage of using the back up and restore method to expand a RAID array?

    Time. The process requires a huge amount of time to implement. However, the method is available to anyone who needs to expand a RAID array, regardless of the operating system or RAID controller being used.

Chapter 12

Review Questions

1. Why might an open slot not be usable?

   It could be that the adjoining slot is in use. In some systems with both ISA and PCI slots, the two types share one common slat on the back, so if one slot is in use, the other isn't available. Another reason might be that the power supply is insufficient.

2. What techniques can you use to free additional interrupt request lines so that they can be used by new add-on cards?

   Consolidate add-on cards so that you have fewer of them. For example, get SCSI controllers with multiple buses instead of one per controller. You can also disable unused ports and devices on the computer. For example, you can disable an unused USB controller or serial port, freeing its interrupt request line.

3. What issues can be associated with I/O ports?

   More than 65,000 I/O ports are available, so running out isn't an issue. However, problems can occur when more than one device tries to use the same I/O port. Fortunately, this is rare in today's environment because of the use of Plug and Play.

4. What information does a compatibility list convey?

   A compatibility list indicates only that the operating system vendor or its agents have tested the device to certify that it works with the operating system. It does not indicate whether the hardware will work with any other hardware. In the case of Linux, the hardware compatibility list may indicate only that there are reports from users of the device working or not working.

Chapter 13

Review Questions

1. What types of batteries are used in most UPSs?

   Most UPSs use sealed lead-acid batteries. They are similar to the ones used in your car, except that they are always sealed, and never require maintenance. Like the battery in your car, however, they have a fixed useful life.

2. What factor do you need to consider when adding additional batteries to a UPS that you don't need to worry about when simply replacing batteries that have reached the end of their life span?

   When adding new batteries, you're changing the potential total output of the batteries. For that reason, you need to confirm that the cables connecting the batteries to the UPS are sufficiently large to conduct all of the current.

3. Why is it advisable to shut down the servers attached to a UPS when replacing the batteries?

   The first reason to shut everything down is to limit the amount of power that is flowing through the UPS, thus limiting the amount of power that you can come into contact with. The second reason is that you may accidentally disconnect power during the operation, and it's better to have shut down the servers normally.

4. What is the line-conditioner mode of a UPS?

   UPSs have transformers and MOVs designed to condition the power most of the time. Line-conditioner mode limits the UPS to using the transformers and MOVs, and thus prevents it from attempting to transfer the load to the inverter and batteries. This mode is used when you know that the batteries or inverter are incapable of supporting the load.

5. What advantages does hardware-based server monitoring have over software-based server monitoring?

   Hardware-based monitoring can continue to function and report even after the motherboard or operating system has failed. The on-board battery backup and separate network connection allow hardware-based monitoring cards to report errors that software-based systems can't report.

6. What conditions in a server's environment do most environmental monitoring systems monitor?

   Environmental monitoring systems generally monitor power, temperature, and noise. They also generally have contacts that allow them to monitor up to three other devices.

7. What are the four ways that software-based monitoring systems can notify you of a problem?

   Console alert, e-mail, numeric pager, and alphanumeric pager. Each has its own advantages, but no matter what your environment is there is a solution to help you become aware of problems on the network.

8. What are the three functions within SNMP?

   SNMP management consoles can read values from an SNMP agent, set values on an SNMP agent, and receive trap messages from SNMP agents. These three functions allow SNMP to perform a variety of functions.

9. What are the disadvantages of polling?

   Polling incurs overhead. The more values you poll and the more frequently you poll for them, the more overhead you create on both the device doing the polling and the device being polled. Polling also means that there is a delay between when

an event occurs and when it is discovered. This delay is inversely related to polling frequency.

10. What are the disadvantages of event-driven or alert-driven notification?

Event-driven notification requires greater complexity on the remote end, but its primary disadvantage is the potential for the event to get lost.

11. How does DMI differ from SNMP?

SNMP is designed for the management of network devices. DMI is designed to manage PCs. SNMP devices are generally fixed, while PCs change with new additions and removals. Because of this, DMI uses a three-tier architecture that uses separate component-based modules that plug into an agent and a third management layer.

12. What function does an MIB serve?

An MIB, or management information base, is a text-based file that identifies all of the objects and values that an SNMP agent supports. It's the roadmap to the SNMP agent.

13. What problems can a multimeter help you diagnose?

A multimeter can help you diagnose problems with the power supply in a server and can also verify end-to-end connectivity within a cable.

14. What is ARP and what does it do?

Address Resolution Protocol (ARP) resolves TCP/IP addresses to hardware addresses. The ARP command allows you to view and control the cache of addresses translated by the ARP protocol.

15. Why would someone use nslookup?

Nslookup is a command to interactively query Domain Name System (DNS) servers. It is used to diagnose problems with the Domain Name System.

16. What other sources, besides checking the hardware compatibility list, should you check when installing a new add-on card?

You should check the vendor's knowledge base and FAQs, as well as the Internet newsgroups. As you learned in Chapter 1, Lesson 1, these alternatives to the documentation are invaluable because they can help identify potential installation problems.

Chapter 14

Review Questions

1. How do incremental and differential backups work for transactional systems?

   Transactional systems maintain transaction logs. These logs are backed up during either a differential or an incremental backup on a transactional system. An incremental backup clears the log once it's been backed up. A differential backup does not clear the transaction log.

2. What is an archive bit, and what does it do?

   The archive bit is an attribute of a file in a FAT or NTFS-based file system that is used to determine whether the file has been backed up. This bit is set automatically whenever the file is modified. The backup software clears the archive bit after it completes a full or incremental backup. Other file systems

support different mechanisms to identify files that may have been backed up.

3. Where did the Tower of Hanoi backup scheme get its name?

   The Tower of Hanoi is a puzzle created by a French mathematician. The tape rotation is based on the same recursive, revolving solution required to solve the Tower of Hanoi puzzle.

4. What is a continuous backup?

   A continuous backup is a special type of backup that backs up the system after every file change or completed transaction. It's used most often in environments where not even a single transaction can be lost. It's been largely replaced by clustered systems.

5. What is tape wear?

   The tape medium eventually wears out. Tapes can be run underneath a tape drive read-write head only so many times. When a tape has been passed underneath the read-write heads enough times, it becomes unstable. A tape shouldn't be used after it's passed underneath the heads the rated number of times.

6. What factor during a backup can significantly affect tape life?

   If a tape drive is not able to maintain motion, the continuous starts, stops, and restarts of the tape drive will wear out the tape quicker. The tape gets stressed when the tape drive stops and restarts. For the tape drive to maintain continuous motion, the system must feed it sufficient information. That is, the server should be able to provide the tape drive with enough information to keep the tape moving. This rate varies if

hardware encryption is being used, but it is generally twice the uncompressed data rate.

7. What are important characteristics in the party responsible for monitoring the backups?

   The person responsible for the backups should have a relatively routine schedule that doesn't change. The more stable the schedule, the more likely it is that the person will remember to monitor the backups.

8. What kind of on-site storage should be provided for tape backups?

   A media-rated fireproof safe should be provided for the on-site storage of tapes. Media-rated fireproof safes keep the tapes safe from excessive heat; a regular fireproof safe doesn't offer enough protection against heat.

9. What are the three options for off-site storage?

   Having an employee take tapes home, hiring a professional off-site storage service, and cross-rotating tapes among different organization sites.

10. What is the best way to determine whether a tape drive needs cleaning? What can you do if this isn't an option?

    The best method for determining when a tape drive should be cleaned is to look for the cleaning indicator on the drive itself. If no indicator is present on the drive, you can estimate the number of hours between cleanings from the tape drive documentation and establish a regular rotation based on that information.

11. Why shouldn't the notification for the backups be set up to notify only on failure?

Because if the backup routine gets stuck and doesn't complete, you won't know there is a problem. If you perform notification only on failures, you won't know that the backup completed successfully, only that it didn't end in failure. It may not have completed at all.

Chapter 15

Review Questions

1. Why are physical housekeeping and verification important, even if the proper monitoring systems are in place?

   Not every potential problem can be monitored. In addition, monitoring systems can fail. Nothing is more reliable than a physical review of the systems.

2. Why is it important to verify climate control and the operational status of fans?

   Servers must be maintained within a specific operating environment. When this range isn't maintained, the useful life of the server is reduced. The climate control equipment establishes this environment. The fans ensure that this environment exists inside the server as well as outside it.

3. When should you specifically test remote access?

   When remote access is limited to providing support for the servers and isn't used by the general user population, it should be tested periodically to ensure that it is still functional.

4. Why is it important to check fault-tolerant systems?

   Because fault-tolerant systems such as RAID arrays, redundant power supplies, and redundant network adapters are designed to hide failures from the users, they are often not

noticed by the administrator. The key is for the server administrator to notice the problem before the next failure happens.

5. Why should surge suppressors be checked?

   Surge suppressors have a finite amount of protection they can offer to the connected equipment, based on the MOVs that are at their core. After the MOVs have been exhausted, they should be replaced.

Chapter 16

Review Questions

1. What is Bloom's cognitive domain?

   It's a representation of the levels at which information can be learned. By progressing from the lower levels to the higher levels, Bloom believed, you achieved a deeper understanding of the material.

2. Why is troubleshooting different from most other learning?

   Troubleshooting requires a deeper understanding of the material and the ability to apply your learning toward a solution. It's the exercise of a process using the information you know to solve a problem.

3. What is the first step in the scientific method?

   To understand the problem. Before you can solve the problem, you must understand it in detail. The first step deals with the collection of information to define the problem.

4. Why doesn't the hypothesis being tested have to be right?

The scientific method is an iterative process. It requires only that something be learned in each iteration. In the next iteration of the process, you use that new information in conjunction with the previous information to come up with another hypothesis.

5. Why is localization of a problem so important?

   Localization is important because it helps to contain the possible problems, and thus the possible solutions, at a manageable level.

6. Why is isolation of a problem important?

   It keeps problems from ganging up on you. Finding two problems at the same time is exponentially harder than finding one problem. Isolation ensures that you're trying to solve only one problem.

7. Who should you call if a hardware driver is responsible for locking up the computer?

   The hardware vendor that provided the device for which the driver was written. The operating system vendor can't help with third-party hardware drivers.

Chapter 17

Review Questions

1. What is the first step in the bootup process?

   The CPU executes the instruction at system memory address FFFF0h. This leads to a jump instruction that points to the location of the POST.

2. Where is the POST stored?

In the ROM on the motherboard. The POST routine also executes POST routines contained in the ROM of add-on cards.

3. What is the first thing you should check when SCSI drives are missing from the bootup process?

   If no devices have been changed, you should check the SCSI BIOS to make sure that the drive is being located and that the BIOS itself is being installed in memory. If any bus configuration changes have been made, it's important to check termination and ID settings.

4. Why are boot sector viruses not a significant issue in today's environment?

   Most software installations today are not done from floppy disks. Boot sector infector viruses require floppy disks in order to spread.

5. What is the most prevalent type of virus today (based on the number of different viruses)?

   The most prevalent type of virus is the file infector virus, one that extends or replaces an executable to include itself. Admittedly, script/macro viruses are catching up at an alarming rate.

6. What is the first step in resolving a network connectivity problem?

   Review the physical layer. Make sure that the cables and hubs are connected so that the link indicators on the network cards and on the hubs themselves are lit. If the physical layer isn't working, nothing will work.

7. Once you've verified the physical layer of connectivity, what steps should you use to diagnose a TCP/IP problem?

Test the local network, test a remote network, and finally test name resolution.

8. Why is DHCP used?

   DHCP simplifies the task of maintaining a TCP/IP network. The DHCP server can give out not only IP addresses but also other TCP/IP configuration information as well.

9. What is a core dump?

   It's an operating system response to a critical error that couldn't be logged in the system log. Core dumps can be either dumps of user information space caused by an error that does not generally affect overall system performance, or system dumps, which are generally produced following errors that result in the system being rebooted, affecting everyone using the system.

10. Why do hard disks vibrate?

    Because the disk is set slightly off center and is spinning at a high rate of speed.